

HEWLETT-PACKARD COMPANY

Intellectual Property Administration

P. O. Box 272400

Fort Collins, Colorado 80527-2400

PATENT APPLICATION
ATTORNEY DOCKET NO. 10992795
IN THE U.S. PATENT AND TRADEMARK OFFICE
Patent Application Transmittal Letter
COMMISSIONER FOR PATENTS

Washington, D.C. 20231

Sir:

 Transmitted herewith for filing under 37 CFR 1.53(b) is a(n): ☒ Utility () Design

☒ original patent application,

☐ continuation-in-part application

INVENTOR(S): Alan H. Karp and Thomas Rokicki
TITLE: FLEXIBLE ALLOCATION OF A RESOURCE

Enclosed are:

☒ The Declaration and Power of Attorney. ☒ signed () unsigned or partially signed

☒ 3 sheets of drawings (one set) () Associate Power of Attorney

☐ Form PTO-1449 () Information Disclosure Statement and Form PTO-1449

☐ Priority document(s) () (Other) _____ (fee \$ _____)

CLAIMS AS FILED BY OTHER THAN A SMALL ENTITY

(1) FOR	(2) NUMBER FILED	(3) NUMBER EXTRA	(4) RATE	(5) TOTALS
TOTAL CLAIMS	17 — 20	0	X \$18	\$ 0
INDEPENDENT CLAIMS	2 — 3	0	X \$80	\$ 0
ANY MULTIPLE DEPENDENT CLAIMS	0		\$270	\$ 0
BASIC FEE: Design (\$320.00); Utility (\$710.00)				\$ 710
TOTAL FILING FEE				\$ 710
OTHER FEES				\$
TOTAL CHARGES TO DEPOSIT ACCOUNT				\$ 710

Charge \$ 710 to Deposit Account 08-2025. At any time during the pendency of this application, please charge any fees required or credit any over payment to Deposit Account 08-2025 pursuant to 37 CFR 1.25. Additionally please charge any fees to Deposit Account 08-2025 under 37 CFR 1.16, 1.17, 1.19, 1.20 and 1.21. A duplicate copy of this sheet is enclosed.

 "Express Mail" label no. EK884232842US

 Date of Deposit 10-13-2000

I hereby certify that this is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to: Commissioner for Patents, Washington, D.C. 20231.

 By Paul H. Horstmann

Typed Name: Paul H. Horstmann

Respectfully submitted,

Alan H. Karp and Thomas Rokicki

 By Paul H. Horstmann
Paul H. Horstmann

Attorney/Agent for Applicant(s)

 Reg. No. 36,167

 Date: 10-13-2-2000

 Telephone No.: (415) 602-1721

UNITED STATES PATENT APPLICATION FOR

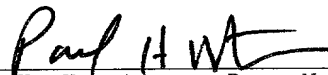
FLEXIBLE ALLOCATION OF
A RESOURCE

Inventors:
Alan H. Karp
Tom Rokicki

CERTIFICATE OF MAILING BY "EXPRESS MAIL"
UNDER 37 C.F.R. § 1.10

"Express Mail" mailing label number: **EK884232842US**
Date of Mailing: **10-13-00**

I hereby certify that this correspondence is being deposited with the United States Postal Service, utilizing the "Express Mail Post Office to Addressee" service addressed to **Assistant Commissioner for Patents, Washington, D.C. 20231** and mailed on the above Date of Mailing with the above "Express Mail" mailing label number.



Paul H. Horstmann, Reg. No. 36,167
Signature Date: 10-13-00

BACKGROUND OF THE INVENTION

Field of Invention

5 The present invention pertains to the field of
computer systems. More particularly, this invention
relates to flexible allocation of a resource in a
computer system.

Art Background

10 A computer system typically includes resources
that are shared among multiple users. An example of
a shared resource is a shared physical memory.
Examples of a shared physical memory include main
memory, persistent memory including mass storage
15 devices, and information stores, etc. Another
example of a shared resource is a communication link.
Yet another example of a shared resource is a
processor.

20 A shared resource usually has a limited capacity
or limited capability with respect to the needs of
the potential users of the shared resource. For
example, a physical memory usually has a limited
storage capacity. A communication link typically has
25 a limited bandwidth. A processor usually has a
limited instruction execution throughput. As a
consequence, computer systems commonly implement
methods for allocating the capacity or capability of
a shared resource among the users of the shared
30 resource.

One prior method for allocating a shared
resource is to employ static partitioning. For

5

10

15

20

SUMMARY OF THE INVENTION

A method is disclosed for flexible allocation of a resource. The method involves assigning a soft limit and a hard limit to each of a set of potential users of the resource. The soft limits are selected to guarantee access to the resource by all of the potential users. The hard limits are selected to enable each potential user to exceed the corresponding soft limit on a first-come-first-served basis. A request from a user for allocation of a portion of the resource is handled by granting the request if the request if allowed would not exceed soft limit assigned to the user. The request is denied if the request if allowed would exceed the hard limit assigned to the user. To avoid overtaxing the capacity of the resource, the request is denied even when the hard limit of the user is not exceeded if the request if allowed would cause a total allocation of the resource to exceed a high watermark assigned to the resource.

Other features and advantages of the present invention will be apparent from the detailed description that follows.

BRIEF DESCRIPTION OF THE DRAWINGS

5 The present invention is described with respect
to particular exemplary embodiments thereof and
reference is accordingly made to the drawings in
which:

10 **Figure 1** shows a computer system that
incorporates the present teachings;

Figure 2 illustrates the handling of a request
for allocation of a resource by a resource manager in
a normal mode in one embodiment;

15 **Figure 3** illustrates the handling of a request
for allocation of the resource by the resource
manager in a reduction mode in one embodiment.

DETAILED DESCRIPTION

5 **Figure 1** shows a computer system 100 that
incorporates the present teachings. The computer
system 100 includes a resource 10 that is shared
among a set of tasks 20-30. Portions of the resource
10 are allocated to the tasks 20-30 by a resource
manager 12. The resource manager 12 maintains a set
of resource allocation parameters 14 which are used
10 in resource allocation.

15 The resource 10 represents any resource having a
limited capacity or capability that may be allocated
among the tasks 20-30. The resource 10 may be a
hardware resource, a software resource, or a
combination hardware/software resource. Examples for
the resource 10 include physical memory such as main
memory, mass storage, persistent stores, information
stores including databases, non-volatile memory,
20 processor time, communication links, and input/output
devices to name a few examples.

25 The tasks 20-30 represent software tasks that
may be executed on the computer system 100. Examples
for the tasks 20-30 include application programs and
related software components and user interface tasks.
Each task 20-30 may be associated with a particular
user of the computer system 10. More than one of the
tasks 20-30 may be associated with the same user. In
30 one embodiment, the resource manager 12 allocates the
resource 10 on a per user basis so that all of the
tasks associated with a given user are confined to a

portion of the resource 10 that is allocated to the given user by the resource manager 12.

5 The computer system 100 may be a single processor system, a multiple processor system, multiple networked computer systems, multiple networked devices which include computing capabilities, or any combination of these. The resource manager 12 may be part of an operating
10 system of the computer system 100, may be a component such as a device driver, and/or may function as a server for the resource 10 that handles requests from the tasks 20-30 which function as clients.

15 The capacity or capability of the resource 10 may be expressed in terms of units. For example, if the resource 10 is a memory then a unit may be a byte, a block, a line, a kilobyte, a megabyte, etc. In another example, if the resource 10 is a
20 communication link then a unit may be a bit per second, a kilobit per second, or a megabit per second of communication bandwidth, etc. In yet another example, if the resource 10 is a processor then a unit may be a million instructions per second (MIPS)
25 of processor execution time.

30 The resource manager 12 receives requests from the tasks 20-30 for allocation of the resource 10. The resource manager 120 allocates portions of the resource 10 to the requesting tasks 20-30 using information provided by the resource allocation parameters 14.

5 The resource allocation parameters 14 include a total capacity or capability (T) of the resource 10 expressed in units. The value of T depends on the characteristics of the resource 10 and the selected units. For example, if the resource 10 is a 1000 megabyte memory then T equals 1000 if the units are megabytes.

10 The value of T may also take into account a portion of the resource 10 which is allocated to system functions and not available to the tasks 20-30. For example, if the resource 10 is a 1000 megabyte memory, then 50 megabytes may be reserved for system use and unavailable for allocation to the tasks 20-30. This yields a value of T of 950 units in megabytes.

20 The resource allocation parameters 14 include a soft limit (S) which applies to each potential user of the resource 10. The soft limit S is a minimum portion of the resource 10 to which each potential user has guaranteed access, thereby preventing potential users from being locked out of the resource 10 at any time.

25 The soft limit S is a tunable parameter of the computer system 100. It is preferable that S be set to a high enough value as to enable advantageous use of the resource 10 but not so high as to needlessly tie up the capacity of the resource 10 when only a few of the potential users access the resource 10.

The soft limit S may be the same for all potential users or may be set on a per user basis or on the basis of classes of users. For example, some classes of users such as those who pay more or those in management positions, etc., may have a higher soft limit than that of ordinary users.

The resource allocation parameters 14 include a hard limit (H) which enables users to exceed their soft limits under predetermined conditions. A given user is always granted his soft limit and may be granted up to his hard limit if the current utilization of the resource 10 can accommodate the request. The maximum value for the hard limit H is equal to T minus the sum of the soft limits of all potential users. The hard limit H is a tunable parameter of the computer system 100. The hard limit H may be the same for all potential users or may be set on a per user basis or on the basis of classes of users.

The resource allocation parameters 14 include a high watermark and a low watermark. The high watermark is an upper limit on the total utilization of the resource 10. The difference between the high and low watermarks provides hysteresis that prevents thrashing that would otherwise occur when one of the tasks 20-30 frees a portion of the resource 10 and then reallocates that portion when the resource 10 is near its capacity.

Figure 2 illustrates the handling of a request 200 for allocation of the resource 10 by the resource

manager 12 in a normal mode in one embodiment. The normal mode of handling a request for allocation is the initial mode before the high watermark of the resource 10 has been exceeded.

5

In this example, the request 200 is generated by the task 20 and specifies a requested portion of the resource 10 expressed as $n1$ units. The request 200 may be an initial request for $n1$ units of the resource 10 or a subsequent request for additional allocation of $n1$ units of the resource 10.

At step 100, the resource manager 12 determines the total allocation of the resource 10 to the user associated with the task 20 that would result if the request 200 is granted. The resource manager 12 records allocations of the resource 10 to users on a per user basis. For example, assume that the task 20 corresponds to user A and that the tasks 21-22 also correspond to user A and have previously been granted $n2$, and $n3$ units of the resource 10, respectively. If so, the total allocation for the user A determined at step 100 is equal to $n2+n3+n1$. If tasks corresponding to the user A have not previously been granted any units of the resource 10 then the total allocation for the user A determined at step 100 is equal to $n1$.

At step 102, the resource manager 12 determines whether the total allocation obtained at step 100 exceeds the soft limit for the user associated with the task 20. If the total allocation obtained at step 100, which includes the request 200 for $n1$

units, would not exceed the user's soft limit then the request 200 is granted at step 104. Otherwise, the user's hard limit is tested at step 106.

5 At step 106, the resource manager 12 determines whether the total allocation obtained at step 100 exceeds the hard limit for the user associated with the task 20. If the total allocation obtained at step 100, which includes the new request 200 for n1
10 units, would exceed the user's hard limit then the request 200 is denied at step 108. Otherwise, the high watermark is tested at step 110.

 At step 110, the resource manager 12 determines
15 whether the total allocation obtained at step 100 would cause the grand total allocation of the resource 10 to all users to exceed the high watermark of the resource 10. If the granting of the request 200 would not cause the grand total allocation to
20 exceed the high watermark then the request 200 is granted at step 114.

 If the granting of the request 200 would cause the grand total allocation of the resource 10 to
25 exceed the high watermark then at step 112 the request 200 is denied. In addition, at step 116 the resource manager 12 enters a reduction mode for handling requests. In the reduction mode, the resource manager 12 always allows requests the reduce
30 the consumption of the resource 10.

Figure 3 illustrates the handling of a request 220 for allocation of the resource 10 by the resource

manager 12 in the reduction mode in one embodiment.
The reduction mode of handling a request for
allocation provides hysteresis that prevents
thrashing that would otherwise occur when one of the
5 tasks 20-30 frees a portion of the resource 10 and
then reallocates that portion when the resource 10 is
near its capacity.

10 In this example, the request 220 is generated by
the task 30 and specifies a requested portion of the
resource 10 expressed as n10 units. The request 200
may be an initial request for the resource 10 by a
user associated with the task 30 or a subsequent
request for additional allocation of n10 units of the
15 resource 10.

At step 130, the resource manager 12 determines
the total allocation of the resource 10 to the user
associated with the task 30 that would result if the
20 request 220 is granted.

At step 132, the resource manager 12 determines
whether the total allocation obtained at step 130
exceeds the soft limit for the user associated with
25 the task 30. If the total allocation obtained at
step 130, which includes the request 220 for n10
units, would not exceed the user's soft limit then
the request 220 is granted at step 134. Otherwise,
the hard limit is tested at step 136.

30 At step 136, the resource manager 12 determines
whether the total allocation obtained at step 130
exceeds the hard limit for the user associated with

the task 30. If the total allocation obtained at
step 130, which includes the request 220 for n10
units, would exceed the user's hard limit then the
request 220 is denied at step 138. Otherwise, the
5 low watermark is tested at step 140.

At step 140, the resource manager 12 determines
whether the total allocation of the resource 10 is
below its low watermark. If the total allocation is
10 not below the low watermark then the request 220 is
denied at step 146.

If the total allocation is below the low
watermark then the request 220 is granted at step
15 142. In addition, at step 144 the resource manager
12 returns to the normal mode for handling requests.

The foregoing detailed description of the
present invention is provided for the purposes of
20 illustration and is not intended to be exhaustive or
to limit the invention to the precise embodiment
disclosed. Accordingly, the scope of the present
invention is defined by the appended claims.

CLAIMS

What is claimed is:

5 1. A method for flexible allocation of a resource,
comprising the steps of:

obtaining a request for allocation of a portion
of the resource;

10 granting the request if the request if allowed
would not exceed a soft limit associated with a
potential user associated with the request, wherein
the soft limit guarantees access to the resource by
the potential user;

15 denying the request if the request if allowed
would exceed a hard limit associated with the
potential user, wherein the hard limit enables the
potential user to exceed the soft limit on a first-
come-first-served basis;

20 denying the request if the request if allowed
would cause a grand total allocation of the resource
to exceed a high watermark assigned to the resource
and granting the request otherwise.

25 2. The method of claim 1, wherein the step of
denying the request if the request if allowed would
cause a grand total allocation of the resource to
exceed a high watermark further comprises the step of
entering a reduction mode for handling a subsequent
request for allocation of the resource.

30 3. The method of claim 2, wherein the reduction
mode comprises the step of:

5 denying the subsequent request if the subsequent
request if allowed would exceed a hard limit
associated with the potential user associated with
the subsequent request;

10

15

20

20

25

30

9. The method of claim 1, further comprising the step of assigning the high watermark to the resource.

10. The method of claim 1, further comprising the
5 step of allocating a portion of the resource for system use.

11. A computer system, comprising:
resource;

10 a set of resource allocation parameters for the resource including a high watermark for the resource and a hard limit and a soft limit associated with a potential user of the resource;

15 task that generates a request for allocation of a portion of the resource;

resource manager that in a normal mode grants the request if the request if allowed would not exceed the soft limit and denies the request if the request if allowed would exceed the hard limit and
20 denies the request if the request if allowed would cause a grand total allocation of the resource to exceed the high watermark and grants the request otherwise.

25 12. The computer system of claim 11, wherein the resource manager switches to a reduction mode if the request if allowed would cause the grand total allocation to exceed the high watermark such that the resource manager grants all subsequent requests that
30 reduce a consumption of the resource while in the reduction mode.

13. The computer system of claim 11, wherein the soft limit is assigned to the potential user to guarantee access to the resource by the potential user.

5

14. The computer system of claim 11, wherein the hard limit is assigned to the potential user to enable the potential user to exceed the soft limit on a first-come-first-served basis.

10

15. The computer system of claim 11, wherein the resource manager enters a reduction mode for handling a subsequent request for allocation of the resource if the request if allowed would exceed the high watermark.

15

16. The computer system of claim 15, wherein the resource manager in the reduction mode grants the subsequent request if the subsequent request if allowed would not exceed a soft limit associated with a potential user associated with the subsequent request and denies the subsequent request if the subsequent request if allowed would exceed a hard limit associated with the potential user associated with the subsequent request and denies the subsequent request if the grand total allocation of the resource is above a low watermark associated with the resource and grants the subsequent request the grand total allocation is below the low watermark.

20

25

30

17. The computer system of claim 16, wherein the resource manager switches to the normal mode if the grand total allocation is below the low watermark.

5
10
15

Attorney Docket No. 10992795

```
graph TD; TM20[TASK 20] <--> RM12[RESOURCE MANAGER 12]; TM21[TASK 21] <--> RM12; TM22[TASK 22] <--> RM12; TM23[TASK 23] <--> RM12; TM24[TASK 24] <--> RM12; TM30[TASK 30] <--> RM12; RAP14[RESOURCE ALLOCATION PARAMETERS 14] <--> RM12; RM12 <--> R10[RESOURCE 10];
```

The diagram illustrates a resource allocation system 100. At the center is a box labeled "RESOURCE MANAGER 12". To its left is a box labeled "RESOURCE ALLOCATION PARAMETERS 14", connected to the manager by a double-headed arrow. Above the manager are six task boxes: "TASK 20", "TASK 21", "TASK 22", "TASK 23", "TASK 24", and "TASK 30", with an ellipsis between "TASK 24" and "TASK 30". Each task box is connected to the resource manager by a double-headed arrow. Below the resource manager is a large box labeled "RESOURCE 10", connected to the manager by a double-headed arrow. The entire system is labeled "100" to the right of the resource manager box.

10992795

**DECLARATION AND POWER OF ATTORNEY
FOR PATENT APPLICATION**ATTORNEY DOCKET NO. 10992795

As a below named inventor, I hereby declare that:

My residence/post office address and citizenship are as stated below next to my name;

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled:

FLEXIBLE ALLOCATION OF A RESOURCE

the specification of which is attached hereto unless the following box is checked:

() was filed on _____ as US Application Serial No. or PCT International Application Number _____ and was amended on _____ (if applicable).

I hereby state that I have reviewed and understood the contents of the above-identified specification, including the claims, as amended by any amendment(s) referred to above. I acknowledge the duty to disclose all information which is material to patentability as defined in 37 CFR 1.56.

Foreign Application(s) and/or Claim of Foreign Priority

I hereby claim foreign priority benefits under Title 35, United States Code Section 119 of any foreign application(s) for patent or inventor(s) certificate listed below and have also identified below any foreign application for patent or inventor(s) certificate having a filing date before that of the application on which priority is claimed:

COUNTRY	APPLICATION NUMBER	DATE FILED	PRIORITY CLAIMED UNDER 35 U.S.C. 119
			YES: _____ NO: _____
			YES: _____ NO: _____

Provisional Application

I hereby claim the benefit under Title 35, United States Code Section 119(e) of any United States provisional application(s) listed below:

APPLICATION SERIAL NUMBER	FILING DATE

U. S. Priority Claim

I hereby claim the benefit under Title 35, United States Code, Section 120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code Section 112, I acknowledge the duty to disclose material information as defined in Title 37, Code of Federal Regulations, Section 1.56(a) which occurred between the filing date of the prior application and the national or PCT international filing date of this application:

APPLICATION SERIAL NUMBER	FILING DATE	STATUS (patented/pending/abandoned)

POWER OF ATTORNEY:

As a named inventor, I hereby appoint the following attorney(s) and/or agent(s) to prosecute this application and transact all business in the Patent and Trademark Office connected therewith:

Customer Number 022879Place Customer
Number Bar Code
Label here

Send Correspondence to:
HEWLETT-PACKARD COMPANY
 Intellectual Property Administration
 P.O. Box 272400
 Fort Collins, Colorado 80527-2400

Direct Telephone Calls To:

Paul H. Horstmann
 (415) 602-1721

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Full Name of Inventor: Alan H. KarpCitizenship: U.S.Residence: 837 Ilima Ct. Palo Alto California 94306Post Office Address: 837 Ilima Ct. Palo Alto California 94306

Inventor's Signature

Alan H. Karp

Date

4 October 2000

ATTORNEY DOCKET NO. 10992795

11 October 2000

Inventor's Signature _____ **Date** _____

Inventor's Signature _____ **Date** _____

Inventor's Signature _____ **Date** _____

Inventor's Signature _____ **Date** _____

Inventor's Signature _____ Date _____

Inventor's Signature _____ Date _____